Predictive Modeling for Natural Gas Prices in Romania Using Time Series Data and Average Temperature

Artemis AIDONI *
Petroleum and Gas University of Ploiesti, Ploiesti, Romania
**Corresponding author, artemisaidoni60@gmail.com*

Konstantinos KOFIDIS[1]
Bucharest University of Economic Studies, Bucharest, Romania
kofidiskostas@hotmail.com

Cristian URSE[2]
Petroleum and Gas University of Ploiesti, Ploiesti, Romania
Cristianurse3@gmail.com

Alexandru STOICA[3]
Petroleum and Gas University of Ploiesti, Ploiesti, Romania
Alexstoica97@yahoo.com

Abstract. *The correct prediction of natural gas prices is an essential tool for energy market decision strategies. The research on deep learning models shows that LSTM models outperform traditional econometric models because they effectively detect the non linear patterns and price instabilities in commodities. The literature review demonstrates that incorporating external variables including weather and economic indicators may lead to better forecasting results. This research implements an LSTM forecasting model to predict Romanian day-ahead natural gas prices by including daily average temperatures as external variables. This study employed partial autocorrelation analysis to determine appropriate lag lengths and RobustScaler to handle outliers for stable model training. The research methodology employs daily data from 2019 to 2024 spitted into training and testing data. The model's performance was evaluated with Mean Squared Error (MSE), Coefficient of Determination (R^2), and Mean Absolute Percentage Error (MAPE). The results showcase a high predictive accuracy, with an MSE of 0.002839, an R^2 of 0.8886, and a MAPE of 4.10% on the test data. These results shows how LSTM with exogenous variables effectively predict market fluctuations when combined together and how temperature acts as a key factor for short and mid range market prediction. The energy sector benefit substantially from accurate forecasts because they enable better risk management, strategic planning and regulatory policy decision-making. The study enhances recent data-driven energy market forecasting research through its implementation of advanced neural architectures and external factors which establishes a base for future investigations into hybrid methods and additional exogenous variables that represent changing market conditions.*

Keywords: Natural Gas, Price Forecasting, Machine Learning, Long Short-Term Memory, Volatility, Temperature, Performance metrics.

Introduction

Natural gas price forecasting plays an important role in modern energy markets, influencing investment strategies, operational planning, and regulatory decisions. As global energy demand continues to grow and diversify, the volatility of natural gas prices creates both opportunities and

DOI: 10.2478/picbe-2025-0194

risks for market participants. Precisely predicting such prices is therefore essential for effective risk management and informed policymaking. Old-style econometric models often struggle to capture the non-linearities and complex interactions inherent in service marketplaces. As a result, modern analytical methods, mainly machine learning and deep learning methods have gathered growing attention.

Research consistently highlights the efficacy of progressive neural network architectures in undertaking forecasting tasks involving noisy and high-dimensional data. Specially, Long Short-Term Memory (LSTM) networks have showed notable success in capturing both instant and long-term temporal dependencies in time-series data. Their gated structure, coupled with advanced regularization techniques, allows LSTMs to address the vanishing and exploding gradient problems that usually hinder simpler recurrent neural networks. Studies of LSTM-based models have pointed bigger performance relative to old-style statistical methods, particularly under conditions of high market volatility.

Despite important progress, numerous challenges remain. First the inherent dynamics of natural gas prices categorized by high, sometimes random variations demand robust architectures capable of learning difficult patterns without surrendering to overfitting. Second, the integration of outside features such as weather, policy alterations, and market sentiment needs model frameworks flexible to include various data sources. Furthermore, balancing computational efficiency with model complexity is vital when applying forecasting tools in real world trading.

Driven by these trials and the proven effectiveness of deep learning, this research uses an LSTM-based methods to forecast daily natural gas day-ahead prices from the Romanian Commodities Exchange. By using robust data scaling, intentionally selected lagged features, and temperature variables, the model aims to understand both short and mid range dependencies which causes price movements. Various evaluation metrics are used to assess predictive performance comprehensively. By doing that this research not only proves the viability of advanced neural models in capturing market trends but also offers understandings into the predictive usefulness of hybrid inputs and data preprocessing techniques.

Literature review

Forecasting natural gas prices is a vital task for energy planning, trading and controlling purposes. Machine learning and deep learning methods improved the ability to predict those prices, addressing challenges created by their volatility and non linear behavior. This review investigates different models applied to natural gas price forecasting, using studies on general commodity forecasting and specific approaches for natural gas markets.

The use of deep learning techniques for forecasting commodity prices has increased substantial in recent years. Many studies highlighted the effectiveness of Long Short-Term Memory (LSTM) networks in seizing temporal dependencies within commodity indices, emphasizing their superior performance in volatile markets. A study used the Bloomberg Commodity Index to assess LSTM models and confirmed their advantage in forecasting commodity price compared to other methods (Ben Ameur et al., 2024). Also, deep learning frameworks have shown significant promise in day-ahead electricity price forecasting. Research in this area concluded that these models consistently outperform traditional statistical approaches, making them highly adaptable for energy markets, including natural gas (Lago et al., 2018).

Moreover, in the context of natural gas price forecasting, different machine learning methods have been applied to address challenges associated with non-linearity and market volatility. Another research studied models such as Artificial Neural Networks (ANN), Support

Vector Machines (SVM), Gradient Boosting Machines (GBM), and Gaussian Process Regression (GPR) for predicting Henry Hub prices, concluding that ANN models were the most effective because of their ability to capture complex, non-linear relations in the data (Su et al., 2019). Furthermore, another study introduced Deep Neural Network (DNN) models that uses fully connected layers and Rectified Linear Unit (ReLU) functions, having greater performance compared to other machine learning techniques (Aliyuda et al., 2021). These findings showcase the importance of advanced models in understanding and predicting complex price patterns.

Additionally, hybrid approaches have also demonstrated excellent potential in improving natural gas price forecasting. For instance, a CEEMDAN–Bagging–HHO–SVR model successfully reduced noise and improved trend capture for monthly Henry Hub prices, showcasing its capability to address market fluctuations (Duan et al., 2023). A similar method is using variational mode decomposition-based methods for interval forecasting, emphasizing their reliability (Ping et al., 2021). Likewise, models including seasonality changes and support vector regression have been utilized to improve the accuracy of natural gas price predictions (Ceperic et al., 2017).

Moving to another hybrid system such as neuro fuzzy model which offers pathways to improve forecast accuracy. An example is the combination of Artificial Neural Networks (ANN), Fuzzy Linear Regression (FLR), and Conventional Regression (CR) that have been successfully applied and improved the estimation of natural gas prices. These models demonstrate their ability to handle noise and non-linearity, showcasing their flexibility across various contexts, including domestic and manufacturing sectors (Azadeh et al., 2012).

Deep learning models mixing Convolutional Neural Networks (CNN) and LSTM layers have shined in capturing short-term and long-term dependencies in U.S. natural gas prices. The CNN layers filter input data to extract essential patterns while reducing noise, and the LSTM layers process these features to recognize temporal dependencies, making the model more effective for chaotic time-series data (Livieris et al., 2020). These methods establish how joining multiple architectures can improve forecast accuracy. Enchaining that previous method, a Temporal Convolutional Network (TCN) model was introduced which uses dilated causal convolutions and dynamic learning rates. This method further refined the ability to predict spot prices with greater precision (Pei et al., 2023). Such advancements underscore the continuous evolution of predictive models in addressing market-specific complexities.

Moreover, market specific factors play a significant role in the selection and success of forecasting models. Like for instance the case of interval forecasts for the Dutch Title Transfer Facility (TTF) which revealed that simpler models like ARMA and Multilayer Perceptrons (MLP) performed better during periods of shocks (Schlueter et al., 2024). Another study on day-ahead price prediction utilized bagged trees and SVM models, demonstrating the importance of incorporating external factors such as stock indices and interest rates into forecasting frameworks (Mouchtaris et al., 2021). These findings highlight that the context of the market often dictates the best suited model.

By using residuals as features in predictive modeling has showed to significantly improves forecast accuracy, particularly during periods of intense market uncertainty such as the COVID-19 pandemic (Jana & Ghosh, 2021).

Also, by including qualitative data has also emerged as a valuable improvement for forecasting models. For example, the addition of news sentiment analysis into predictive frameworks has revealed to improve accuracy by complementing quantitative data with qualitative insights (Tang et al., 2019). In the same way, research on hedging and speculative behaviors has showed their roles in influencing market dynamics, offering valuable insights into mechanisms that can stabilize prices and mitigate volatility (Klein, 2017). These methods align with broader efforts to expand the inputs considered in forecasting models, ensuring they are both comprehensive and responsive to real-world dynamics.

Finally, the literature demonstrates the increasing complexity of models used for natural gas price forecasting. From advanced deep learning methods like LSTM and CNN to hybrid such as CEEMDAN–Bagging–HHO–SVR, these approaches successfully address challenges like non-linearity, volatility, and external influences. Nevertheless, model selection remains highly dependent on content, requiring careful consideration of market conditions, data accessibility, and forecasting approaches. As natural gas markets changes, future research should focus on enhancing model adaptability, integrating diverse data sources, and improving predictive precision to ensure relevance in an increasingly complex and dynamic energy landscape.

Methodology

LSTM

The first model is the LSTM which is widely viewed as the most effective and frequently used method for time series forecasting. This model is mainly used for tackling complex problems due to their ability to retain and utilize short-term and long-term data patterns (Livieris et al., 2020; Aidoni et al., 2023). A Long Short-Term Memory (LSTM) cell is a specialized type of Recurrent Neural Network (RNN) that explodes the vanishing gradient issues often seen in standard RNNs. Each LSTM cell maintains two states:

- Hidden state h_t
- Cell state c_t

At each time step t, the cell receives an input vector xt (the features at time t), as well as the previous hidden state $h\ t-1$ and cell state ct-1. The new hidden and cell states ($h\ t$,) are computed using three gates:

- Forget Gate f_t
- Input Gate i_t
- Output Gate o_t

The operations are described by:

Input gate:

$$i_t = \sigma(b_i + U_i \cdot x_t + W_i \cdot y_{t-1}) \quad (1)$$

Forget gate:

$$f_t = \sigma\left(b_f + U_f \cdot x_t + W_f \cdot y_{t-1}\right) \quad (2)$$

Output (Hidden) State:

$$y_t = o_t \circ tanh(c_t) \quad (3)$$

Output Gate:

$$o_t = \sigma(b_o + U_o \cdot x_t + W_o \cdot y_{t-1}) \quad (4)$$

Candidate Cell State:

$$\tilde{c}_t = tanh(b_c + U_c \cdot x_t + W_c \cdot y_{t-1}) \quad (5)$$

Cell State:

$$c_t = f_t \circ c_{t-1} + i_t \circ \tilde{c}_t \quad (6)$$

An important hyperparameter in an LSTM time series model is the sequence length (SEQ_LENGTH), which specifies how many past time steps the model observes before making a prediction. In this study, we integrate the Non-linear Auto-Regressive with Exogenous Inputs (NARX) model as a foundation for selecting this sequence length and incorporating exogenous variables effectively. Using the NARX framework ensures that the delay variable is grounded in statistical significance, enhancing the model's predictive capacity. The NARX model is defined as:

$$\hat{Y}_{(t+p)} = f(Y_t^{(d_Y)}, X_t^{(d_X)}) \quad (7)$$

where:

$Y_t^{(d_Y)} = (Y_t, Y_{t-1}, \dots, Y_{t-d_y+1})$ represents the lagged values of the target variable Y, up to d_Y delays,

$X_t^{(d_X)} = (X_1^{d_{X_1}}, X_2^{d_{X_2}}, \dots, X_N^{d_{X_N}})_t$ is the set of lagged values of the N exogenous input variables, $X_i^{(d_{X_i})} = (X_{i,t}, X_{i,t-1}, \dots, X_{i,t-d_{X_i}+1})$, for each exogenous input variable $X_{i,} i = 1,2, \dots, N$

here d_y and d_{X_i}represent the respective delays for the target variable Y and the exogenous variables X_i. The combined delay variable d is chosen as d=max(d_Y, d_{X_1},…, d_{X_N}), capturing the most relevant temporal dependencies across both target and exogenous variables.

When p=1, the NARX model corresponds to a one-step-ahead prediction. In our work, we use a simplified version of the model with a single delay variable, $delay$, expressed as:

$\hat{Y}_{t+1} = f(Y_t, Y_{t-1}, \dots, Y_{t-delay+1}, X_{1,t}, X_{1,t-1}, \dots,$
$X_{1,t-delay+1}, \dots, X_{N,t}, X_{N,t-1}, \dots, X_{N,t-delay+1})$ (8) (Cocianu & Uscatu, 2025)

In our implementation, SEQ_LENGTH is set to 4, meaning that at each training step, the LSTM receives a block of four consecutive days' data—encompassing lagged price and temperature features—as input. The are two reasons behind choosing four days: (1) it provides enough recent historical context to capture short- to mid-range dependencies, and (2) it avoids introducing unnecessary noise and computational overhead that can arise from longer sequences.

From Figure 1, we observe that partial autocorrelation values for price remain notable up to approximately the fourth lag, suggesting that information from the preceding four days is especially relevant.

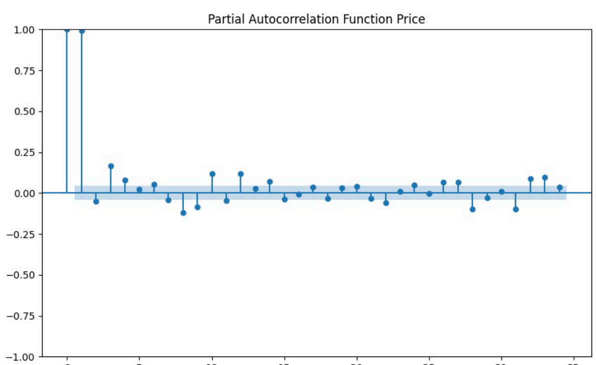

Figure 1. Partial Autocorrelation Function for Price

Source: Authors' own research results.

Similarly, the partial autocorrelation for average temperature also indicates significant influence up to the fourth lag, reinforcing our decision that four days of temperature history is sufficient to capture meaningful patterns.

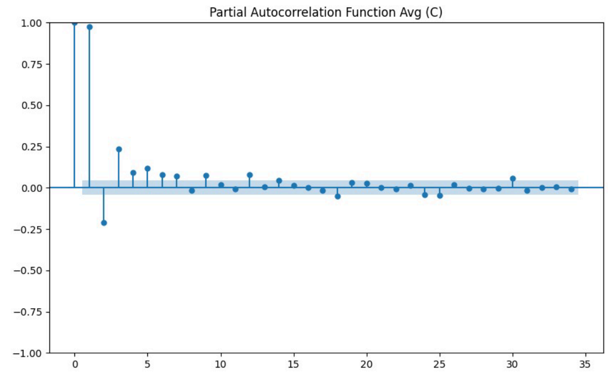

Figure 2. Partial Autocorrelation Function for Avg temperature (°C)

Source: Authors' own research results.

Empirically, using a 4-day window (i.e., SEQ_LENGTH = 4) has proven effective in balancing the need for capturing relevant temporal structure while avoiding the drawbacks of an excessively long sequence.

In time series forecasting, incorporating lagged features enables models to capture temporal dependencies by utilizing historical data. This approach allows the model to understand past behavior over a specified number of periods (Koutlis et al., 2020).

To capture these dependencies, we created lagged price features, helping the model see historical price behavior over a chosen number of days. We constructed 1 lagged column, where price$t$ is the price at day t, as shown:

$$price_{lagk_t} = price_{t-k}\ , k = 1 \qquad (9)$$

The number of hidden units in the LSTM model depends on the dimension of the input layer, which is the number of considered time series multiplied by the delay value.

We remove rows with missing values that arose from generating these lagged features. After preparing the data in this way, we then design our LSTM-based model as follows:

- LSTM (64 units)
    - Activation: tanh
    - Kernel regularizer: L_2(0.001)
- Dense (32 units)
    - Activation: tanh
    - Kernel regularizer: $L2$(0.001)
- Dense (1 unit)
    - Activation: Linear (no explicit activation function, suitable for regression)

After defining the Enhanced LSTM model, the code executes a series of epochs to iteratively optimize the model parameters:

For epoch=1,…,150:

1. Forward and Backpropagation
    - The model performs forward passes on Xtrain, computing predictions and comparing them with y train.

- Gradients are calculated based on the MSE loss, and the Adam optimizer updates the weights accordingly.

2. Performance Tracking.
   - Metrics such as MSE and MAE are recorded. The training process halts automatically after 150 epochs (no early stopping or hyperparameter search is included in the current code).

Upon completion of these epochs, the model parameters are fixed at their final, converged state.

Performance metrics

In this study, we employ a comprehensive set of evaluation metrics to assess the accuracy and robustness of our forecasting model. These metrics include:

- Mean Absolute Percentage Error (MAPE), one of the most widely used measures of forecast accuracy (Kim & Kim, 2016),
- Coefficient of Determination (R^2), used to measure the linear correlation between the actual and the forecasting values (Mashaly et al., 2015) and
- Mean Squared Error (MSE), generally amenable to more in-depth mathematical or statistical analyses (Willmott, 1982).

The above-mentioned evaluation indices are defined as follows:

$$MSE = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2 \qquad (10)$$

$$MAPE = \frac{1}{N}\sum_{i=1}^{N}\frac{|y_i - \hat{y}_i|}{y_i} \cdot 100 \qquad (11)$$

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{N}(\bar{y} - \hat{y}_i)^2} \qquad (12)$$

where y_i is the actual value and $\hat{y}_i$ is the model's forecasting value, N is the total number of samples and $\bar{y}$ is the mean of the real values.

Results and discussions

Dataset description

The dataset utilized in this study contains daily observations of natural gas day-ahead prices (in LEI/MWh) from the Romanian Commodities Exchange (Bursa Română de Mărfuri, BRM) and average daily temperatures for Romania in Celsius from meteorological sources. The data covers the period from January 3, 2019, to October 22, 2024, resulting in 2,109 entries. The dataset provides a comprehensive representation of market trends and climatic variability over nearly six years.

To facilitate robust model training and evaluation, the dataset was divided into two subsets:

- Training set: This subset includes data from January 3, 2019, to August 27, 2023, spanning 1,687 daily entries. The training set covers a diverse range of market conditions and ensures that both long-term and short-term trends are represented. This comprehensive temporal range allows models to learn intricate patterns in the data.
- Testing set: This subset consists of data from August 28, 2023, to October 22, 2024, spanning 422 daily entries. The days without recorded trades resulted in missing values. Since these days do not provide price data, they were dropped from the dataset to ensure the model is trained

only on valid trading days. This period provides a sufficient amount of unseen, out-of-sample data to rigorously evaluate the performance of forecasting models and ensure their generalizability.

A detailed statistical summary of the two key variables, natural gas prices, and temperatures, is provided in Table 1.

Table 1. Descriptive Statistics of Natural Gas Prices and Average Daily Temperatures

Metric	Gas Price (LEI/MWh)	Temperature (°C)
Mean	216.62	13.18
Std. Dev.	204.71	8.56
Minimum	31.61	-8.3
Median	139.99	13.1
Maximum	1,207.87	30.6
Skewness	1.93	-0.09
Kurtosis	3.72	-1.09

Source: Authors' own research results.

Data scaling

Different features in a dataset can vary widely in magnitude. To facilitate stable training of neural networks, RobustScaler is employed to scale both features (X) and target (y) to more consistent ranges. The RobustScaler is particularly effective for data containing outliers because it uses the interquartile range instead of the global min/max (Qian et al., 2022).

Let x be a raw feature, median (x) be the median of x, and IQR (x) the interquartile range of x. Robust scaling transforms each x to x' as:

$$x' = \frac{x - median(x)}{IQR(x)} \qquad (13)$$

which makes the data less sensitive to outliers than other methods such as Min-Max scaling. Outliers were handled using the method mentioned above, which normalizes data by removing the median and scaling according to the interquartile range (Qian et al., 2022).

Results

The model was trained using the following configuration:

- Train/Validation Split: 80% of the data was allocated for training, with 20% reserved for validation to monitor and mitigate overfitting.
- Hyperparameters:
  - Epochs: 150
  - Batch Size: 256

To ensure robustness, the model was executed 10 times, and the average performance metrics were computed.

The model's performance was evaluated using a range of metrics to assess its predictive accuracy across both the test dataset and the entire dataset. Table 2 provides a detailed summary:

Table 2. Evaluation metrics results

Metric	Test Data	Full Data
Mean Squared Error (MSE)	0.002839	0.047288
R-squared (R²)	0.8886	0.969
Mean Absolute Percentage Error (MAPE)	4.10%	6.05%

Source: Authors' own research results.

The metrics highlight the model's robust predictive performance, as evidenced by the low MSE and MAPE values across both datasets. The high R^2 values further confirm the model's ability to explain the variance in natural gas prices effectively.

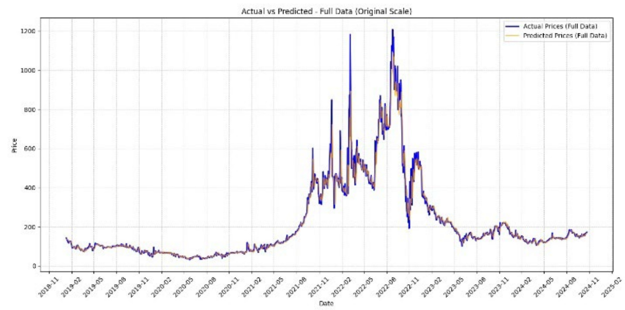

Figure 3. Actual vs. Predicted – Full Dataset

Source: Authors' own research results.

Figure 3 illustrates the comparison between actual and predicted prices across the full dataset. The model demonstrates its capability to capture both long-term trends and short-term fluctuations, with minor deviations during periods of heightened volatility.

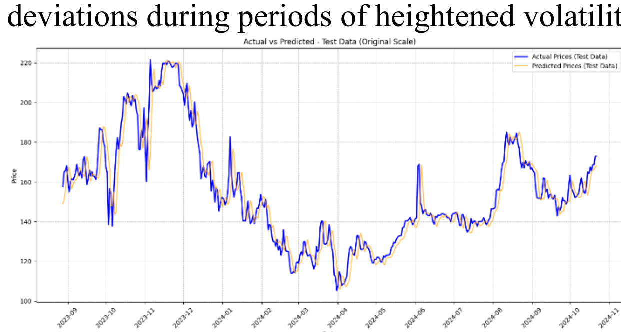

Figure 4. Actual vs. Predicted – Test Dataset

Source: Authors' own research results.

Figure 4 provides a detailed view of the model's performance on the test dataset. The alignment between actual and predicted values suggests the model's strong generalization ability.

Figure 5 illustrates the scatter plot of predicted versus actual prices for the full dataset, providing a visual representation of the model's accuracy. Each point in the plot corresponds to an individual prediction, with actual prices on the x-axis and predicted prices on the y-axis. The red dashed diagonal line represents the ideal fit, where a perfect predictive model would place all points directly along this line. The distribution of points along the diagonal suggests a strong correlation between predicted and actual values, indicating that the LSTM model successfully captures the underlying price trends. For lower and mid-range price levels, the model demonstrates particularly strong predictive performance, with minimal deviations from actual values. However, in higher price ranges, some degree of dispersion is observed, suggesting that while the model effectively

DOI: 10.2478/picbe-2025-0194, pp. 2505-2515, ISSN 2558-9652 |
Proceedings of the 19th International Conference on Business Excellence 2025

tracks price movements, it encounters greater difficulty in precisely capturing extreme fluctuations. This behavior is expected, as price spikes often result from sudden market events or external factors that may not be fully accounted for in historical patterns.

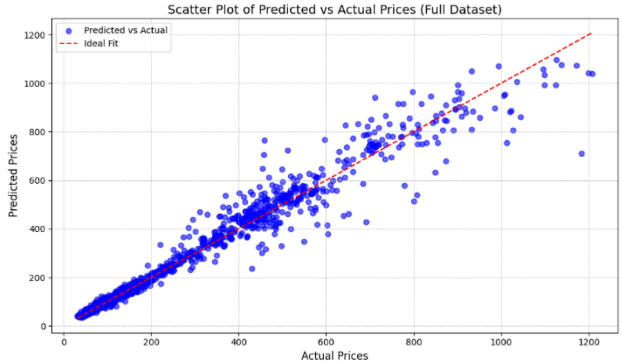

Figure 5. Scatter Plot: Predicted vs. Actual

Source: Authors' own research results.

Conclusion

This study explores the field of natural gas price forecasting by demonstrating the effectiveness of an LSTM-based model trained on a dataset containing many years of Romanian day-ahead prices and daily temperatures. Adding the temperature as an exogenous variable and with the use of partial autocorrelation to determine a proper lag length, showed to boost predictive performance. By using a RobustScaler to manage outliers and mitigate scale discrepancies, the model remains stable and generalizes well to unseen data. The assessment metrics, mainly the high R^2 and low MAPE scores, validates the models ability to understand both long and short term market variations.

In conclusion, despite these promising results, some limitations exist and require further investigation. We need to investigate the use of more extensive hyperparameter tuning and advanced regularization strategies that could optimize LSTM performance further. Also, adding more exogenous factors such as global economic indicators or geopolitical events can enhance the model's predictive accuracy. Finally, exploring hybrid techniques that integrate convolutional layers, attention mechanisms, or fuzzy logic may create improvements in scenarios where market shocks demand greater flexibility. Future research dedicated to these modifications can help refine forecasting models, thus offer more robust and actionable insights for stakeholders in the energy sector.

References

Ben Ameur, H., Boubaker, S., Ftiti, Z., Louhichi, W., and Tissaoui, K. (2024). Forecasting commodity prices: empirical evidence using deep learning tools. *Annals of Operations Research*, 339(1), pp.349–367.

Lago, J., De Ridder, F., and De Schutter, B. (2018). Forecasting spot electricity prices: Deep learning approaches and empirical comparison of traditional algorithms. *Applied Energy*, 221(1), pp.386–405.

Su, M., Zhang, Z., Zhu, Y., Zha, D., and Wen, W. (2019). Data Driven Natural Gas Spot Price Prediction Models Using Machine Learning Methods. *Energies*, 12(9), p.1680.

Aliyuda, A., Ahmed, M.K., Aliyuda, K., and Bello, A.M. (2021*). Deep Neural Network Model for Improving Price Prediction of Natural Gas*. International Conference on Data Analytics for Business and Industry (ICDABI), Sakheer, Bahrain, pp.113-117.

Duan, Y., Zhang, J., and Wang, X. (2023). Henry Hub monthly natural gas price forecasting using CEEMDAN–Bagging–HHO–SVR. *Frontiers in Energy Research*, 11, art. 1323073.

Livieris, I.E., Pintelas, E., Kiriakidou, N., and Stavroyiannis, S. (2020). An Advanced Deep Learning Model for Short-Term Forecasting U.S. Natural Gas Price and Movement. *IFIP Advances in Information and Communication Technology*, 585, pp.165–176.

Pei, Y., Huang, C.J., Shen, Y., and Wang, M. (2023). A Novel Model for Spot Price Forecast of Natural Gas Based on Temporal Convolutional Network. *Energies*, 16(5), p.2321.

Schlueter, S., Pappert, S., and Neumann, M. (2024). Interval Forecasts for Gas Prices in the Face of Structural Breaks – Statistical Models vs. Neural Networks. *arXiv*:2407.16723.

Mouchtaris, D., Sofianos, E., Gogas, P., and Papadimitriou, T. (2021). Forecasting Natural Gas Spot Prices with Machine Learning. *Energies*, 14(18), p.5782.

Jana, R.K., and Ghosh, I. (2021). A residual driven ensemble machine learning approach for forecasting natural gas prices: analyses for pre-and during-COVID-19 phases. *Annals of Operations Research*, 301(1), pp.123–143.

Ceperic, E., Žiković, S., and Ceperic, V. (2017). Short-Term Forecasting of Natural Gas Prices using Machine Learning and Feature Selection Algorithms. *Energy*, 140(1), pp.893–900.

Ping, J., Zhenkun, L., Jianzhou, W., and Lifang, Z. (2021). Decomposition-selection-ensemble forecasting system for energy futures price forecasting based on multi-objective version of chaos game optimization algorithm. *Resources Policy*, 73, p.102234.

Tang, Y., Wang, Q., Xu, W., Wang, M., and Wang, Z. (2019). *Natural Gas Price Prediction with Big Data*. IEEE International Conference on Big Data, Los Angeles, USA, pp.5326–5330.

Azadeh, A., Sheikhalishahi, M., and Shahmiri, S. (2012). A hybrid neuro-fuzzy simulation approach for improvement of natural gas price forecasting in industrial sectors with vague indicators. *Int. J. Adv. Manuf. Technol*., 62, pp.15–33.

Klein, T. (2017). Dynamic correlation of precious metals and flight-to-quality in developed markets. *Finance Research Letters*, 23, pp.283–290.

Kim, S., and Kim, H. (2016). A new metric of absolute percentage error for intermittent demand forecasts. *International Journal of Forecasting*, 32(3), pp.669-679.

Willmott, C.J. (1982). Some comments on the evaluation of model performance. *Bulletin of the American Meteorological Society*, 63(11), pp.1309–1313.

Mashaly, A.F., Alazba, A.A., Al-Awaadh, A.M., and Mattar, M.A. (2015). Predictive model for assessing and optimizing solar still performance using artificial neural network under hyper arid environment. *Solar Energy*, 118, pp.41-58.

Cocianu, C.-L., and Uscatu, C.R. (2025). Towards a New MI-Driven Methodology for Predicting the Prices of Cryptocurrencies. *Electronics*, 14(1).

Koutlis, C., Papadopoulos, S., Schinas, M., and Kompatsiaris, I. (2020). LAVARNET: Neural Network Modeling of Causal Variable Relationships for Multivariate Time Series Forecasting. *Applied Soft Computing*, 96.

Qian, H., Wen, Q., Sun, L., Gu, J., Niu, Q., and Tang, Z. (2022). *RobustScaler: QoS-Aware Autoscaling for Complex Workloads*. IEEE 38th International Conference on Data Engineering (ICDE), pp.2762–2775.

Aidoni, A., Kofidis, K., Cocianu, C.L., and Avram, L. (2023). Deep learning models for natural gas demand forecasting: a comparative study of MLP, CNN and LSTM. *Romanian Journal of Petroleum & Gas Technology,* 4(1), pp.133-148.